Germs, Genes, and Memes:
Function and Fitness Dynamics on Information Networks

Patrick Grim,[1,4] Daniel J. Singer,[2] Christopher Reade,[3] and Steven Fisher[4]

[1]Department of Philosophy, Stony Brook University, Stony Brook, NY 11794
[2]Department of Philosophy, University of Pennsylvania, Philadelphia, PA 19104
[3]Gerald R. Ford School of Public Policy, University of Michigan, Ann Arbor, MI 48109
[4]Center for Study of Complex Systems, University of Michigan, Ann Arbor, MI 48109
singerd@phil.upenn.edu, pgrim@notes.cc.sunysb.edu

Abstract

It is widely accepted that the way information transfers across networks depends importantly on the structure of the network. Here, we show that the *mechanism* of information transfer is crucial: in many respects the effect of the specific transfer mechanism swamps network effects. Results are demonstrated in terms of three different types of transfer mechanism: germs, genes, and memes. With an emphasis on the specific case of transfer between sub-networks, we explore both the dynamics of each of these across networks and a measure of their comparative fitness.

Germ and meme transfer exhibit very different dynamics across linked networks. For germs, measured in terms of time to total infection, network type rather than degree of linkage between sub-networks is the primary factor. For memes or belief transfer, measured in terms of time to consensus, it is the opposite: degree of linkage trumps network type in importance. The dynamics of genetic information transfer is unlike either germs or memes. Transfer of genetic information is robust across network differences to which both germs and memes prove sensitive.

We also consider function: how well germ, gene, and meme transfer mechanisms can meet their respective objectives of infecting the population, mixing and transferring genetic information, and spreading a message. A shared formal measure of fitness is introduced for purposes of comparison, again with an emphasis on linked sub-networks.   Meme transfer proves superior to transfer by genetic reproduction on that measure, with both memes and genes superior to infection dynamics across all networks types. What kinds of network structure optimize fitness also differ among the three. Both germs and genes show fairly stable fitness with added links between sub-networks, but genes show greater sensitivity to the structure of sub-networks at issue. Belief transfer, in contrast to the other two, shows a clear decline in fitness with increasingly connected networks.

When it comes to understanding how information moves on networks, our results indicate that questions of information dynamics on networks cannot be answered in terms of networks alone. A primary role is played by the specific mechanism of information transfer at issue. We must first ask about how a particular type of information moves.

Information Transfer Mechanisms: Dynamics and Fitness

Questions about how information is transferred on networks are of importance in many disciplines. Biologists seek to understand how signals are transferred across animal groups (Strandburg-Peshkin, Twomey, Bode, Kao, Katz, Ioannou, Rosenthal, Torney, Wu, Levin & Couzin 2013, Handegard, Boswell, Ioannou, Leblanc, Tjøstheim & Couzin 2012). They also ask how genetic information is transferred between individuals and populations in sexual and reproductive networks (Chowdhury, Lloyd-Price, Smolander, Baici, Hughes, Yli-Harja, Chua, & Reibeiro, 2010; Holland 2012). Computer scientists ask how to optimally transfer information across computer networks and between networks of sensors (Wooldridge 2002, Tanenbaum 2006, Science Publishing Group, 2013-).   Most public health interventions are information-based; public health professionals need to know the optimal ways of targeting and delivering that information (Hawe & Potkin 2009; Grim, Reade, Singer, Fisher, & Majewicz, 2010a, 2010b; Grim, Thomas, Fisher, Reade, Singer, Garza, Fryer & Chatman 2012). Marketing executives want to know how information spreads by word of mouth and over social networks and the effect that has on acceptance of innovation and product adoption (Valente 1995; Achrol & Kotler 1999; Gordon 1999; Easley & Kleinberg 2010). Policy makers are interested in information flow across social communication networks with an eye to policy formulation, promulgation, and enforcement (Atkinson & Coleman 1992; Peterson 2003).

Questions about the dynamics of information are also of importance in philosophy, most clearly in philosophy of information (Floridi 2003, 2011; Adriaans & van Bnethem, 2008) but in other areas as well. Social epistemologists and philosophers of science ask how the transfer of information and the social structure of science affect the knowledge of the community (Grim 2009; Grim, Singer,

Fisher, Bramson, Berger, Reade, Flocken & Sales 2013; Zollman 2010, 2013; Alexander 2013). Questions of ethics and the evolution of norms can turn on the dynamics of information as well (Alexander & Skyrms 1999; Bendor & Swistak 2001; Alexander 2007).

Genes, germs, and memes all carry information. Viruses and bacteria contain genetic information in RNA and DNA. The genes of higher organisms carry information in sexually transferred DNA. Propositional information is represented in memes and beliefs. Each of these forms of information has a characteristic mechanism of transfer. Germs use the asexual reproduction characteristic of infection, genes of higher organisms use sexual reproduction with crossover, and beliefs are transferred through learning, reinforcement, and accommodation. In the first part of this paper, we use computer-aided simulations to explore the dynamics of these three forms of information on linked sub-networks.

We model information transfer with information pooling models, in the sense of Zollman 2013. In each round of the computer simulation, information for each agent or node of the network is affected only by the information at points to which it is connected. It is well established for such models that network structure affects the amount of time required until a community—until information spread across the network is roughly uniform (e.g. Zollman 2007). The main result of part I of the paper is that information transfer *type* is crucial to issues of time convergence. If the timing of spread across a community is the concern, whether because we want to delay the spread of viral infection or we want to optimize scientific consensus, it is crucial that we first ask how the particular kind of information at issue moves.

Germs, genes, and memes differ in having a distinctive mechanism of transfer, but each of these types of information also serves certain ends. A highly fit pathogen is one that infects a host population so as to optimize reproduction and spread. A fit gene is one that under selective pressure best survives and spreads across a population. A fit meme or belief is one that accurately represents the world, optimizes cognitive function, or otherwise benefits its agent (Grim 2009).

In part II, we take some first steps toward exploring differences in fitness between these three types of information transfer, again with an emphasis on linked sub-networks. Our results indicate how different types of information can be used to optimize different goals, but also make clear some of the costs of those optimizations: what germs, genes, and memes are distinctively good at, but their specific limitations as well.

In terms of a formally uniform measure of fitness we introduce, meme transfer produces generally higher results than the other two types of information across all network types. What kinds of network structure optimize fitness for each type, however, depends on the specific mechanism at issue. The fitness of both sexual and asexual reproduction—germs and genes—is fairly resilient to the addition of links between sub-networks, though sexual reproduction shows more sensitivity to the structure of networks that are linked. Belief, in contrast to the other two, shows a decline in fitness with increasingly connected networks.

What our results indicate quite generally is that questions of information dynamics on networks cannot be answered in terms of network structure alone. We must first know the specific method of information transfer at issue. There are specific things that information transfer on the model of germs, genes, and memes is good for, on specific kinds of networks, but always with a specific cost. It is only once we restrict ourselves to a specific mode of transfer that we can speak of consistent effects of network structure per se.

Information Networks

Real social networks aren't uniform and homogenous. Social communities are composed of sub-communities, with varying degrees of contact in terms of the physical contact necessary for disease transmission, the sexual contact necessary for genetic mixing, and the communication streams necessary for belief transfer. For animals, sub-communities divided by geographical and ecological barriers are crucial in understanding both disease transmission and the genetics of speciation. In the case of people, sub-communities are divided along racial, ethnic, demographic, and socio-economic lines. In order to understand infection dynamics we need to understand physical contact networks, including contact links and gaps between sub-communities. In order to understand genetic change we need to understand links of sexual contact between sub-communities. In the case of belief we need to understand the impact of linkages between sub-communities not only of physical contact but of communication and trust.

Figure 1 shows a series of four networks related in terms of their structure. The network on the left is a single total network. The three pairs on the right form paired total sub-networks with increasing numbers of connecting links.
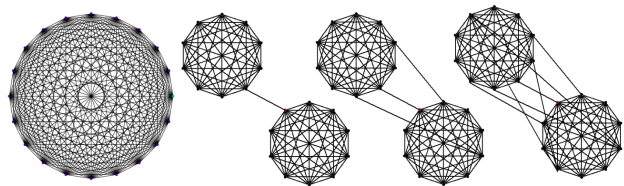

Fig. 1 A single total network and increased degrees of linkage between total sub-networks

We can measure connectedness of sub-networks in terms of the number of linkages between nodes of distinct

groups or sub-networks as a percentage of the total possible. Linkages between sub-networks of this type have also been termed 'bridges,' analogous to a concept of bridges in computer networking and identified in Trotter, Rothenberg and Coyle (1995) as a key area for future work in network studies and health care. L. C. Freeman (1977) speaks of the same phenomena in terms of segregation and integration between sub-networks.

Figure 2 shows the types of linked sub-networks we concentrate in the present study: linked total networks, rings, small worlds, random and scale-free networks. For simplicity we use just two sub-networks of 50 nodes each (Figure 2 uses a smaller number of nodes for visibility). Our rings use just one connection to a single neighbor on each side. For small worlds we work with single rings in which roughly 9% of nodes have been re-wired at random. In our random networks roughly 4.5% of possible connections are in each sub-network. Our scale-free networks are constructed by the preferential attachment algorithm of Barabási and Albert (1999). Where needed, we add a minimal number of links to assure a connected network in each case; hence the 'roughly' of 9% rewired nodes and 4.5% of possible connections in the case of small world and random sub-networks.
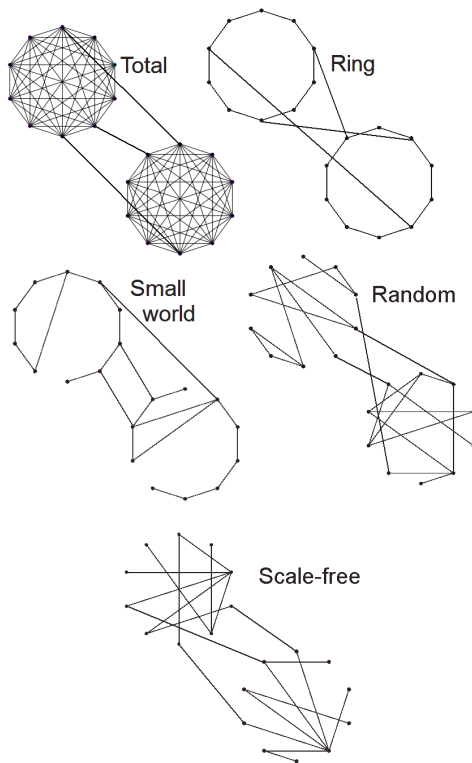

Fig. 2 The network types at issue. Simulation runs used sub-networks of 50 nodes; smaller numbers of nodes are shown here for visibility.

In the course of our investigation we vary both the types of sub-networks and the number of links between them.

An important part of the work also calls for a comparison between results for linked networks of a given type and results for a single 100-node network of that type. We compare results for ring sub-networks with a particular number of links between them and a single ring with the same number of added links. Such a method allows us to focus on the effects of two distinct aspects of network structure:

(a) network type—ring, wheels, hubs, small worlds, random or scale-free, and

(b) degree of linkage between sub-networks.

The idea is to use differences between results on single networks and on linked sub-networks to tease out those aspects of over-all network structure that crucial for particular results. Similar results between a single network of a particular type and linked sub-networks of that type is evidence that it is the *network type* rather than details of linkage that is important for the result. Different results between a single network and linked sub-networks—particularly where differences carry across different networks types—is evidence that it is *degree of linkage* between the sub-networks that is doing the work.

Part I: Network Dynamics for Germs, Genes, and Memes

Infection

Germs offer the simplest case. Here we vary (a) the structure of sub-networks involved—whether sub-networks at issue are rings, small worlds, random, scale-free, or total—and (b) the degree of linkage between those sub-networks. What effect do each of these have on the dynamics of infection?

Regardless of the infection rate, time to total infection turns out to be importantly sensitive to network structure. It is not sensitive, however, to whether that structure is instantiated in a single network or in linked sub-networks. Figure 3 shows results from simulation for increased linkages between sub-networks. For each number between 1 and 50 we create 1000 networks with random links of that number between sub-networks, starting with a single infected node. Figure 3 shows the average steps to total network infection with over the 1000 runs, here using a 100% infection rate. Figure 4 shows results in which links are added not between sub-networks but within a single large network of each structure.
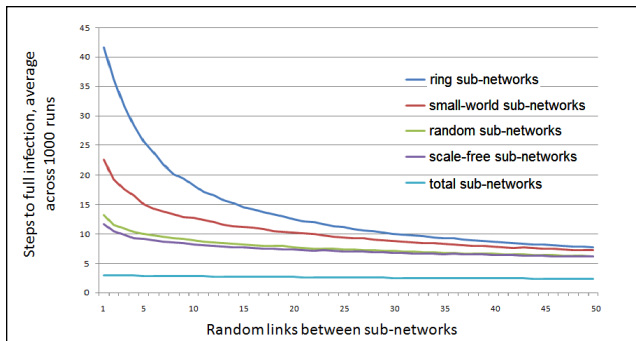
Fig. 3 Average time to total infection with increasing links between ring, small world, random, scale-free, and total sub-networks
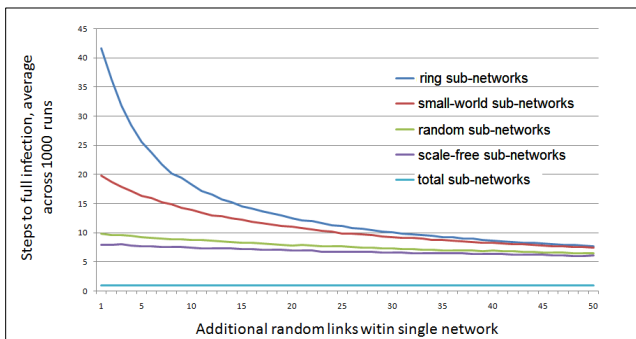

Fig. 4 Average time to total infection with increasing links added in single ring, small world, random, scale-free, and total networks

Results in the two cases are virtually identical. The difference in plotted lines in each figure shows that network structure does make a significant difference in time to total infection, but the fact that such a structure is instantiated in sub-networks rather than a single network does *not* make a difference. In all the cases considered, it is not degree of linkage sub-networks but the network type of either single or linked sub-networks—whether ring, small world, random, scale-free, or total—that produces network-specific signatures for infection.[1]

If one wants to either plot or predict, explain or understand the course of an epidemic—or of information transfer analogous to an epidemic—what one needs to know is first and foremost not the degree of linkage between sub-networks but the characteristic structure of the sub-networks themselves.

Memes and beliefs

Like germs, memes spread across social networks, but in this case the dynamics of information diffusion are dramatically different. Some earlier work has trumpeted similarities in infection dynamics and the spread of ideas

---

[1] This largely accords with analytic results by Golub and Jackson (forthcoming) regarding the role of linkage in diffusion dynamics.

(Newman 2001, Redner 1998, Börner et. al. 2003). Our purpose is to emphasize crucial differences between them.

In our models, agents' beliefs are represented as a single number between 0 and 1, which we can think of an estimate of a fixed quantity. Perhaps these are beliefs in the severity of a disease, the probability of contracting the disease, or the effectiveness of vaccination. (Harrison, Mullen, & Green 1992; Janz & Becker, 1984; Mullen, Hersey, and Iverson, 1987; Strecher & Rosenstock, 1997). Agents are influenced by the beliefs of those around them, updating their beliefs in terms of the beliefs of those with whom they are informationally linked.

Though simple, we regard this model of belief as sufficiently realistic for our purposes. Some beliefs *can* be represented on such a scale, and it is clear that people's beliefs are influenced, among other things, by the beliefs of those with whom they have contact. What is admittedly unrealistic is the simple form of belief updating we use: an averaging of beliefs with those with whom one has network contact (French 1956, Harary 1959, DeGroot 1974, Golub & Jackson 2010). Averaging of beliefs in an informational neighborhood certainly does not capture the full complexity of belief change. What an averaging mechanism does mimic, however, in a conveniently simple way, is the pattern of reinforcement often characteristic of belief change. The more one's beliefs are like those of more of one's network neighbors, the less inclination there will be to change those beliefs. The more one's beliefs are out of sync with one's neighbors, the greater the pressure there will be to change one's beliefs (Asch 1952, 1955, Bond & Smith 1996; Cialdini & Golstein 2004).

Using belief averaging, and regardless of initial assignment of belief values, all agents in this model eventually approach the same belief value. We can therefore measure the effect of network structure on belief convergence by measuring the number of steps required on average until all agents in the network are within, say, a range of .1 above or below the mean belief across the network as a whole. In what follows we use this range of variance from the mean as our measure of convergence, averaging over 100 runs in each case.

We begin with polarized agents. Half of our agents are drawn from a pool with belief measures that form a normal distribution around .25, with a standard deviation of .06. The other half are drawn from a pool with belief measures in similar normal distribution around .75. In the case of single networks agents are drawn randomly from each pool. In studying linked sub-networks our agents in one sub-network are drawn from the .25 pool; those in the other are drawn from the .75 pool. Belief polarization of this form is necessary to study the effects of sub-network linkage in particular; were beliefs of our agents randomized within each sub-network, convergence to an approximate mean could be expected to occur in each sub-

network independently. Time to consensus would not then be a measure of the effect of sub-network linkage.

In outlining the dynamics of infection above, we contrasted linked sub-networks of particular structures—ring, small world, random, total, and scale-free—with single networks of the same structure. In exploring the dynamics of belief we again study these side by side. Figure 5 shows graphs indicating times to belief convergence for each of our network types. Times to consensus with increased linkages between sub-networks of a given type are shown in blue. Times to consensus with increased linkages within single networks of that type are shown in red.
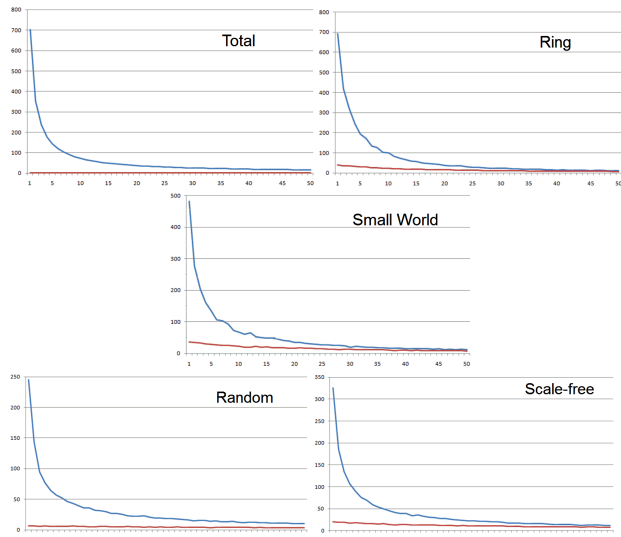
Fig. 5 Times to belief convergence in various networks for increasing links between sub-networks (shown in blue) and within single networks of that type (shown in red).

Notice two important features of these graphs: (1) the extreme divergence between rates of belief convergence between linked sub-networks and single networks in each case, and (2) the remarkable similarity of the curves for linked sub-networks in each case. That similarity is emphasized by plotting results for all sub-network types together in log-log form in Figure 6, revealing the classic signature of a power law.

What these results indicate is that where information is transferred in the manner of memes rather than germs, the dominating effect of increased linkage between sub-networks is independent of the structure of the sub-networks themselves. For meme transmission, unlike infection, the degree of linkage between sub-networks trumps network type. If one wants to trace the course of an epidemic, we noted, it is crucial that one knows the structure of networks involved. To plot or predict, explain or understand the course of belief transmission, in contrast,

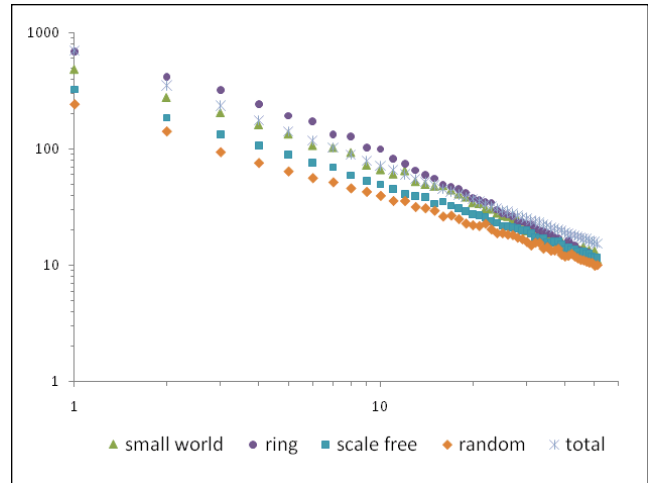it is degree of linkages between sub-networks, of whatever type, that is crucial.

Fig. 6 Log-log plots of times to belief consensus with increased linkages between sub-networks of various types

Genes

Genetic information transfer in higher organisms is characterized by crossover in sexual reproduction.[2] To simulate this form of information transfer, we gave each agent a genetic code consisting of a binary string of length 100. Half of the population starts out with a genetic code of all ones, the other with all zeroes. In the case of linked sub-networks, each sub-network begins with a uniform genetic code of either zeroes or ones. In the case of single networks, we randomize the two codes in the population.

On each time step of the model, each agent pairs off with an unpaired agent she is connected to, if there is such an agent. Each pair then mates and two new genetic codes are formed. Each new code is a crossover of the two parents' codes, code from one parent to the left of a random crossover point and code from the other parent to the right. The two new codes will generally differ because each code is produced with a random crossover point. The two parents are then replaced by offspring bearing the new genetic codes but with the same network connections.

In the limit, in connected networks, we can expect this form of genetic updating to produce a uniform genetic code across the population. As the network converges, genetic information will become more uniformly spread. We can therefore use convergence within a given threshold of similarity as a measure of the dynamics of genetic information. We take as our convergence measure that point at which two agents drawn randomly from the population can be expected to differ in less than 20% of

---

[2] Both asexual organisms and viruses transfer genetic information by reproduction as well. Here we use 'genetic' to refer solely to the crossover characteristic of sexual reproduction.

their genetic code. For our binary strings of length 100, this measure represents a Hamming distance of less than 20.

For belief transfer, we've seen, degree of linkage between sub-networks is the major factor in time to convergence. For infection, the crucial factor is network and sub-network type rather than degree of linkage. Genetic information transfer exhibits a mixture of these features, but it doesn't match either of the other patterns in all respects.
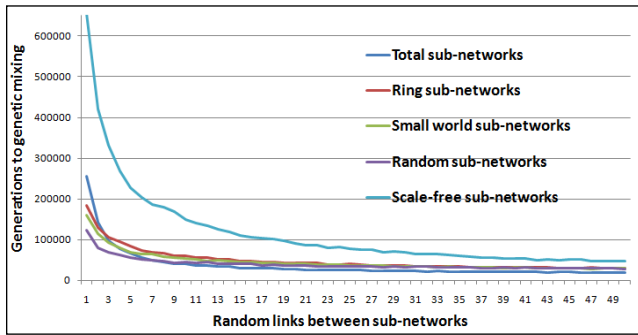

Fig. 7 Generations to genetic convergence with a Hamming distance of 20 for increased linkages between sub-networks of various types.

Figure 7 shows genetic dynamics results for linked sub-networks of each of our types. Here, like in the case of belief, network type tends to make very little difference. The data from our scale-free preferential attachment networks are outliers; we remain unsure why. When that case is removed, the proximity of results for increased linkages and regardless of the types of sub-networks linked is even clearer. Those results are shown in a re-scaled graph in Figure 8 and in log-log form in Figure 9.
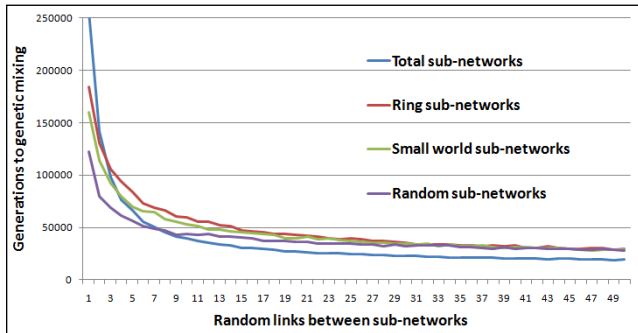

Fig. 8 Generations to genetic convergence with a Hamming distance of 20 for increased linkages between total, ring, small world and random sub-networks.
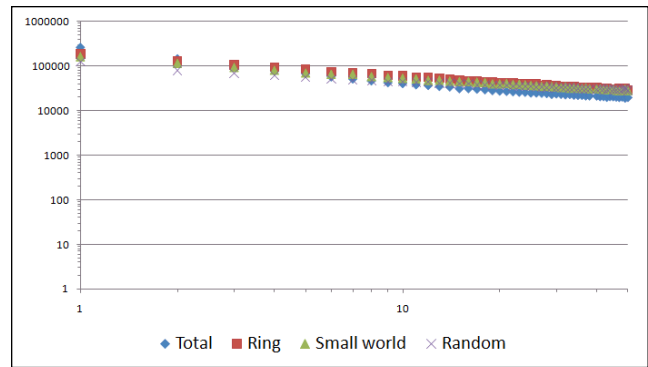

Fig. 9 Log-log plot of genetic convergence for increased links between total, ring, small world and random sub-networks.

Our results indicate that if information is being transferred via a genetic mechanism, the number of links between sub-networks trumps the structure of the sub-networks themselves when it comes to time to convergence. This is a feature genetic information shares with meme or belief transfer. Here as in that case we again have the signature of a power law, though the slope or scaling exponent is very different. Genes therefore share some notable features with memes.

In the case of memes, however, we found a major difference between results for linked sub-networks and for single networks of a given type. That difference did not appear in infection dynamics: there convergence times on linked networks of a given type closely paralleled those for single networks with the same number of added links. In this respect, genetic information transfer turns out to be more like infection. Figure 10 shows comparisons for our graph types between added linkages between networks (shown in blue) and within a single network of the same type (shown in red). Total networks stand out, but that is because 'added' links in a total network are redundant. In all other cases single networks start with a lower time to convergence, but after just a few added links, the times to convergence for single and linked sub-networks are nearly identical. In this respect genes are more like germs.
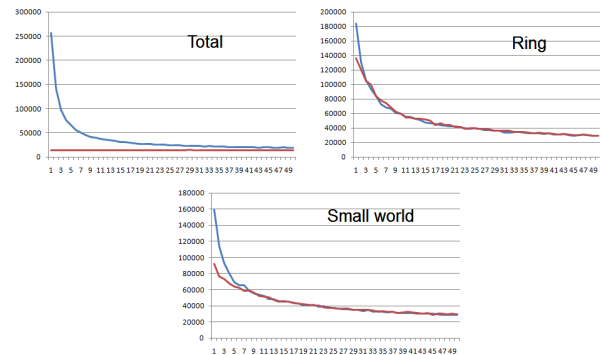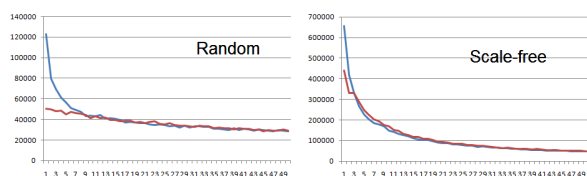
Fig. 10 Times to genetic convergence in various networks for increasing links between sub-networks (shown in blue) and within single networks of that type (shown in red).

Genetic information transfer therefore shows features of both germ and meme dynamics, but is clearly different from either. For memes, the primary determinant of time to convergence is the number of links between sub-networks, not the structure of the linked networks themselves. For germ transfer, the pattern is reversed: network type proves far more important than degree of linkage. For genes, increased network connectivity does increase speed to convergence, but with a dynamics in which neither network type nor degree of linkage play the dominant role characteristic of germs or memes.

Part II: Fitness for Germs, Genes, and Memes

Different forms of information transfer act in the service of different ends. Here, again using linked sub-networks as a primary tool, we take some first steps toward a more rigorous understanding of the comparative fitness of different forms of information transfer toward different ends.

We encode information for asexual and sexual reproduction—germs and genes--as binary strings of length 21. In the previous section, we modeled belief as a real number between 0 and 1, which creates an obstacle for comparison across the three information transfer mechanisms. By encoding real values for belief as binary decimals of length 21, however, we have a consistent way of representing information across all three transfer mechanisms.

Our fitness measure will also be the same across the three mechanisms. We use .100000000000000000000 in binary, or .5 in decimal, as our 'optimal' code: that genetic string most strongly selected for in the environment of germs or genes, or that information string most strongly selected for in an environment of memes. In the case of belief, our optimal code might be thought of as an indicator of truth or action effectiveness. In all cases, we measure fitness by the arithmetic distance between an agent's code and that arbitrary fitness target.

We use the same range of networks types as before, with increased linkages (i) between sub-networks and (ii) within single networks of each type. In each case we start with a randomization of information strings across all

nodes of the network. What differs is the dynamics of information transfer.

For the asexual reproduction of germs, a node is replaced by that node with the highest fitness to which it is linked in the network. In these early studies we do not include any provision for mutation.

For sexual reproduction—genes—the information code of a node is replaced with a cross-over at a random point between it and the code of its fittest neighbor, again without mutation.[3]

For belief dynamics, the real value of each node is averaged with that of its fittest neighbor.

With a uniform method of representing both information and fitness across the three types of information transfer, we can compare fitness dynamics across different types of networks, whether single or linked sub-networks, and with different degrees of linkage. In each case we measure the number of generations required until the network converges; all nodes have values within .001 of the average fitness of nodes across the network at convergence. At that point we can also measure how 'fit' the final convergence is: proximity to the optimal target at convergence. We use 100 runs for each linkage value of each network, deriving both mean and standard deviation across those results.

The first measure of comparison is time to convergence. In this regard the reproduction of germs is clearly the fastest. Updating by meme or belief reinforcement is the second fastest, with genetic crossover a much slower third. Figure 11 offers just one example: relative speed to convergence with additional links added to a single scale-free network. Slower mechanisms, with more generations to fixation, are displayed toward the top. Very similar results hold across all networks in our sample. Germs are faster than memes. Memes are faster than genes.
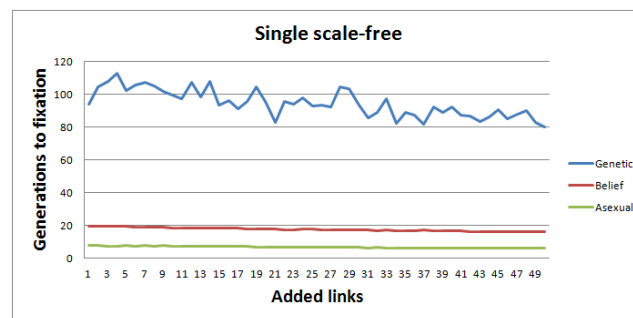

Fig. 11 Relative speed to convergence of the transfer mechanisms with added links in a single scale-free network, typical of relative speed across all networks.

---

[3] In order to avoid bias at crossover in favor of information toward the front of a genetic code, the significance of each position on the string is randomly decided at the beginning of each simulation run.

What of fitness measured as proximity to our optimal code at convergence? Figure 12 shows relative fitness for the three information transfer types, again typical of results for our networks as a whole. The more fit, with closer proximity to our optimal value, are displayed toward the top of the graph

The asexual information transfer of germs is the fastest in terms of time to convergence, but scores worst in terms of fitness. The reinforcement dynamics of memes or belief prove the best. The fitness of genetic recombination lies in the middle. On this measure, memes prove fitter than genes. Genes prove fitter than germs.
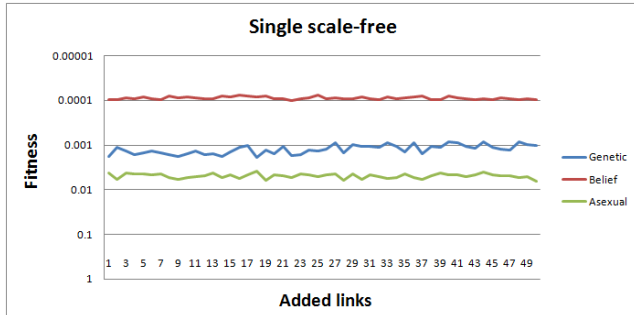
Fig. 12 Log plotting of relative fitness for the information transfer mechanisms

Orderings in terms of speed and fitness clearly differ for our information transfer mechanisms, but those orderings are not simply inverted. The fastest information transfer mechanism is also that with the lowest fitness: the asexual reproduction of memes. But the slowest transfer mechanism is not that with the highest fitness. The slowest is genetic cross-over, but it is the belief reinforcement mechanisms of memes that show the highest fitness.

Across a sample of networks, Zollman (2007, 2010) demonstrates a direct trade-off between speed to convergence and accuracy for networks of agents pooling information about bandit problems. Our work here shows that such a trade-off does not hold across all types of information pooling. Belief transfer proves both faster to convergence and more fit than genetic transfer, for example. This strengthens the lesson from the previous section: to understand the dynamics of information across networks, we must first ask what mechanism of information transfer is under consideration.

There are devils in the details, however, and a more nuanced story is told by the dynamics of different transfer mechanisms on various networks. Figures 13 and 14 show patterns for the asexual reproduction of germs across linked sub-networks of our various types. Speed to convergence varies widely across network type, increasing with increased linkages, but without any sharp differences in fitness.
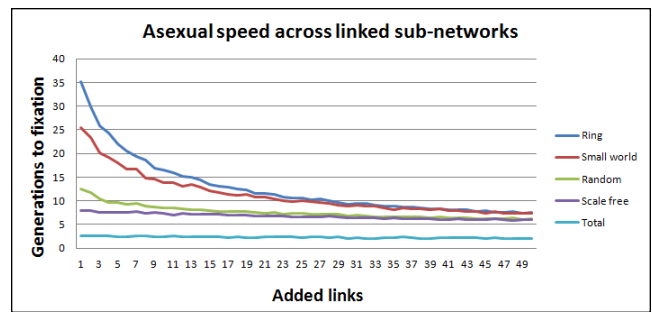

Fig. 13 Speed to asexual convergence with added links between all types of sub-networks
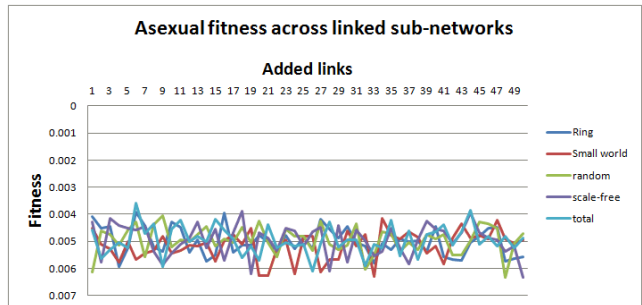

Fig. 14 Asexual fitness at convergence with added links between all types of sub-networks

Figures 15, 16 and 17 show comparative results for memes. Figure 15 makes it clear that linked total networks are clearly the fastest to convergence, with a uniform speed. In the other cases speed increases with added linkages, particularly in the case of relatively distributed networks like our rings and small worlds.
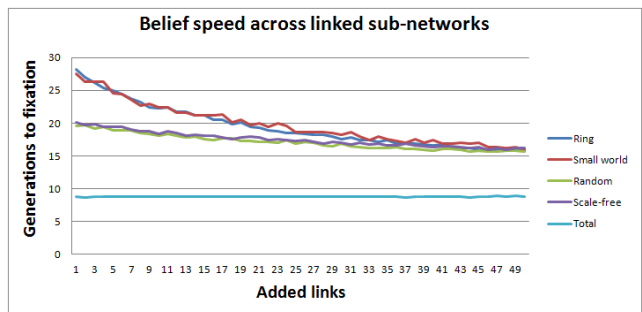

Fig. 15 Speed to belief convergence with added links between all types of sub-networks.

Figure 16 shows that the speed to convergence for total networks is made up for by very poor fitness. Figure 17 shows results for the other networks in a finer resolution, indicating decreasing fitness with increased linkages between sub-networks in all cases except scale-free networks, where the results are unclear. Meme fitness decays with increased linkages between sub-networks. These results are in accord with related work on epistemic networks (Grim 2009; Grim, Reade, Singer, Fisher & Majuewicz 2010). There too increased linkage between

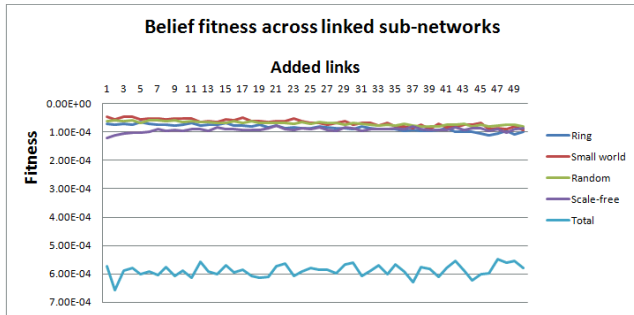sub-networks leads to decay in belief-based information transfer.

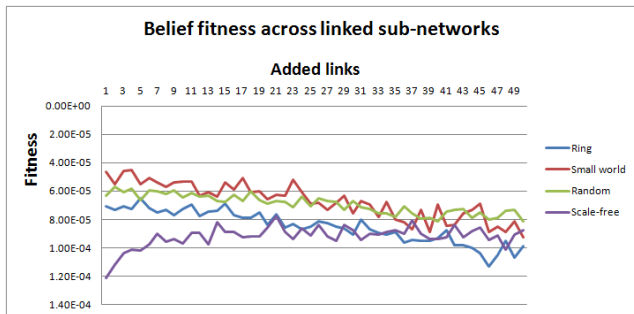Fig. 16 Belief fitness at convergence with added links between all types of sub-networks

Fig. 17 Decay in belief fitness at convergence with added links between all types of sub-networks

Comparisons regarding speed and fitness for genes are shown in Figures 18 and 19. Though significantly slower, genetic information transfer shows the same pattern of speeds to convergence on various network types that the infection transfer of germs do, with a comparable impact of increased linkages and network types. Like germs, but unlike memes, gene fitness also shows relative constancy across added linkages within a network type. Whereas the fitness of germs seems essentially random across fitness types, however, fitness in the case of genetic transfer does appear to be higher for distributed networks such as rings and small worlds, in that respect echoing a characteristic that also appears with memes.
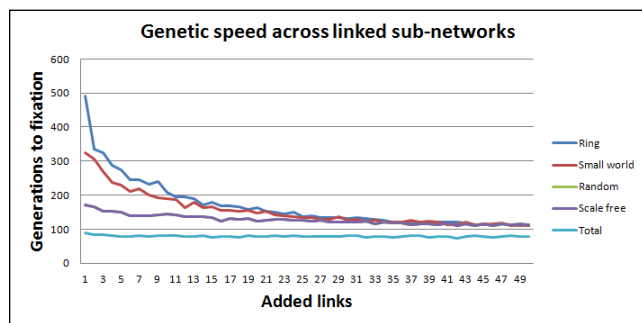

Fig. 18 Speed to genetic convergence with added links between all types of sub-networks
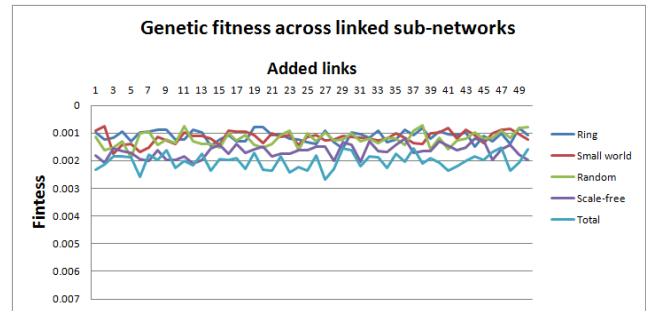

Fig. 19 Genetic fitness at convergence with added links between all types of sub-networks

As a whole, our results indicate that the information transfer by reinforcement characteristic of memes or beliefs is an outlier compared to information transfer by either the asexual reproduction of germs or the sexual reproduction of genes. The information transfer characteristic of memes quickly produces convergence and generates highly fit communities. But it comes with a cost—a marked decrease in fitness with increased network linkages that does not appear with either germs or genes.

Conclusion

How does information move on networks? Our attempt here has been to show that such a question has no answer—no single answer—independent of a specification of the particular mechanisms of information transfer at issue.

Germs, genes, and memes represent different information strategies, with different dynamics in the case of linked sub-networks and intriguingly different sensitivities to degree of linkage and network type. Our attempt has been to take some first steps to comparing fitness regarding those three types as well. Different methods of information transfer can be expected to optimize different kinds of fitness, selected for in different environments and measured in terms of different ends.

The importance of understanding network structure for transfer of information has received a great deal of well-deserved attention. That understanding remains radically incomplete, however, without a similar focus on various modes of information transfer as different as germs, genes, and memes.

Acknowledgements

References

Achrol, R. S., & Kotler, P. (1999). Marketing in the Network Economy. *The Journal of Marketing*, 63, 146-163.

Adriaans, P. w. & J. F. A. K. van Bentham (2008). *Handbook of Philosophy of Information*. Elsevier Science Publishers.

Alexander, Jason McKenzie (2007). *The Structural Evolution of Morality*. Cambridge University Press.

Alexander, Jason Mckenzie, & Brian Skyrms (1999). "Bargaining with Neighbors: Is Justice Contagious?" *Journal of Philosophy* 96, 588-598.

Asch, S. (1952). *Social Psychology*. Englewood Cliffs, N.J.: Prentice-Hall.

Asch, S. (1955). Opinions and Social Pressure. *Scientific American* 193, 31-35

Atkinson, Michael M., & William D. Coleman, 1992. "Policy Networks, Policy Communication and the Problems of Governance," *Governance 5*, 154-180.

Barabási, A.-L. and Albert, R. (1999). Emergence of scaling in random networks. *Science* 286, 509-512.

Bendor, Jonathan & Piotr Swistak (2001). The Evolution of Norms. *American Journal of Sociolo*gy 106, 1493-1545.

Bond, R., & P. B. Smith (1996). "Culture and Conformity: A Meta-Analysis of Studies Using Asch's Line Judgment Task," *Psychological Bulletin* 119, 111-137.

Chowdhury, S., J. Lloyd-Price, O-P. Smolander, W. C. V. Baici, T. R Hughes, O. Yli-Harja, G. Chua, & A. S. Ribeiro, 2010. "Information propagation within the Genetic Network of Saccharomyces cerevisiae," *Systems Biology* 4: 143.

Cialdini, R. B., & N. J. Goldstein (2004). "Social Influence: Compliance and Conformity," *Annual Review of Psychology* 55, 591-621.

DeGroot, M. H. (1974). Reaching a Consensus. *Journal of the American Statistical Association* 69, 118-121.

Easley, David & Jon Kleinberg, *Networks, Crowds and Markets: Reasoning About a Highly Connected World*, 2010. Cambridge Univ. Press.

Floridi, Luciano 2003). Ed., *The Blackwell Guide to the Philosophy of Computing and Information*. Oxford: Blackwell.

Floridi, Luicano (2011). *The Philosophy of Information*. Oxford University Press.

Freeman, L. C. (1978). Segregation in social networks. *Sociological Methods and Research* 6: 411-429.

French, J. (1956). A Formal Theory of Social Power. *Psychological Review* 63, 181-194.

Golub, B. & M. O. Jackson (2010). Naïve Learning in Social Networks: Convergence, Influence, and the Wisdom of Crowds. *American Economic Journal: Microeconomics*, 2, 112-149.

Golub, B., & Jackson, M. O. (forthcoming). How homophily affects learning and diffusion in networks. Available as a working paper at www.bepress.com/feem/paper296

Gordon, A. W. (1999). Network effects in marketing. *Marketing Research*, 11(3), 36

Grim, P. (2009). Threshold phenomena in epistemic networks. *Proceedings, AAAI Fall Symposium on Complex Adaptive Systems and the Threshold Effect*, FS-09-03, AAAI Press.

Grim, P., Reade, C., Singer, D. J., Fisher, S. and Majewicz, S. (2010a). What you believe travels differently: Information and infection dynamics across sub-networks," *Connections* 30, 50-63.

Grim, P., C. Reade, D. J. Singer, S. Fisher, & S. Majewicz, 2010b. "Robustness across the Structure of Sub-Networks: The Contrast between Infection and Information Dynamics," Proceeds, AAAI Symposium on Complex Adaptive Systems: Resilience, Robustness and Evolvability, FS-10-03.

Grim, P., S. B. Thomas, S. Fisher, C. Reade, D. Singer, M. A. Garz, C. S. Fryer & J. Chatman, 2012. "Polarization and Belief Dynamics in the Black and White Communities: An Agent-Based Network Model from the Data, *Artificial Life 13*, MIT Press.

Grim, P., D. J. Singer, S. Fisher, A. Bramson, W. Berger, C. Reade, C. Flocken & A Sales (2013). "Scientific Networks on Data Landscapes: Question Difficulty, Epistemic Success and Convergence." *Episteme* 10, 441-464.

Harary, F. (1959). A Criterion for Unanimity in French's Theory of Social Power. In *Studies in Social Power*, ed. by D. Cartwright. Ann Arbor: University of Michigan Press

Harrison, J. A., P. D. Mullen, & L. W. Green (1992). A meta-analysis of studies of the health belief model. *Health Education Research* 7:107–116.

Hawe, P., & L. Potkin, 2009. "What is population health intervention research?' *Canadian Journal of Public Health* 11(1):Suppl 18-14.

Holland, John, *Signals and Boundaries*, 2012. Cambridge, MA: MIT Press

Janz, N. K., & M. H. Becker (1984). The health belief model: A decade later. *Health Education Quarterly* 11:1–

Mullen, P. D., J. Hersey, & D. C. Iverson (1987). Health behavior models compared. *Social Science and Medicine* 24:973–981.

Peterson, John, Policy Networks, 2003. Wien: Institut für Höhere Studien (IHS), Vienna: Institute for Advanced Studies. http://www.ihs.ac.at/publications/pol/pw_90.pdf

Strandburg-Peshkin, A., Twomey, C.R., Bode, N.W., Kao, A.B., Katz, Y., Ioannou, C.C., Rosenthal, S.B., Torney, C.J., Wu, H., Levin, S.A. & Couzin, I.D. (2013) Visual sensory networks and effective information transfer in animal groups, *Current Biology* 23(17), R709-711.

Science Publishing Group (2013 -). *International Journal of Sensors and Sensor Networks.*

Strecher, V. J., and Rosenstock, I. M. (1997). The health belief model. In . Glanz, F. M. Lewis, and B. K. Rimer, eds., *Health Behavior and Health Education: Theory, Research, and Practice.* San Francisco: Jossey-Bass.

Tanenbaum, Andrew S. (2006). *Computer Networks, 4ᵗʰ Edition*. Pearson Eduction.

Trotter, Robert T. II, Richard B. Rothenberg, Susan Coyle (1995). Drug abuse and HIV prevention research: Expanding paradigms and network contributions to risk reduction. 18(1): 29-45.

Wooldridge, Michael (2002). *An Introduction to MultiAgent Systems*. John Wiley & sons.

Zollman, Kevin J. S. (2007). The communication structure of epistemic communities. *Philosophy of Science* 74 (5):574-587.

Valente, Thomas W. (1995) 1995. *Network Models of the Diffusion of Innovation*. Cresskill, New Jersey: Hampton Press.

Zollman, Kevin J. S. (2010). The epistemic benefit of transient diversity. *Erkenntnis* 72 (1):17 - 35.

Zollman, Kevin J. S. (2013). Network Epistemology: Communication in Epistemic Communities. *Philosophy Compass* 8 (1):15-27.